

# PERFORMANCE IMPROVEMENTS OF PROBABILISTIC TRANSCRIPT-ADAPTED ASR WITH RECURRENT NEURAL NETWORK AND LANGUAGE-SPECIFIC CONSTRAINTS

Xiang Kong<sup>1</sup>, Preethi Jyothi<sup>2</sup>, Mark Hasegawa-Johnson<sup>2</sup>

<sup>1</sup> Department of Computer Science, University of Illinois at Urbana-Champaign, IL, USA

<sup>2</sup> Beckman Institute, University of Illinois at Urbana-Champaign, IL, USA

{xkong12, pjyothi, jhasegaw}@illinois.edu

## ABSTRACT

Mismatched transcriptions have been proposed as a mean to acquire probabilistic transcriptions from non-native speakers of a language. Prior work has demonstrated the value of these transcriptions by successfully adapting cross-lingual ASR systems for different target languages. In this work, we describe two techniques to refine these probabilistic transcriptions: a noisy-channel model of non-native phone misperception is trained using a recurrent neural network, and decoded using minimally-resourced language-dependent pronunciation constraints. Both innovations improve quality of the transcript, and both innovations reduce phone error rate of a trained ASR, by 7% and 9% respectively.

**Index Terms:** G2P, RNN Encoder-Decoder, probabilistic transcription, mismatched crowding, automatic speech recognition system

## 1. INTRODUCTION

In building standard automatic speech recognition (ASR) systems, the first step is to collect sufficient amounts of transcribed speech data. Crowdsourcing [1] [2] has been regarded as a useful resource to speed up this step, by acquiring transcripts from a large number of crowd workers. However, normal crowdsourcing methods fail for under-resourced languages since they require transcribers to be native speakers of the target language, and it is hard to find native speakers of under-resourced languages online. As a matter of fact, according to the survey [3], most languages are under-resourced which makes crowdsourcing only applicable to a small fraction of the languages in the world.

Mismatched crowdsourcing was introduced as a potential way to solve the limitation of crowdsourcing. In mismatched crowdsourcing, workers transcribe a language they do not know at all. Instead of using meaningful words in the target language, they just use nonsense syllables in their own orthography (here: English letters) to transcribe speech data [4]. These so-called mismatched transcripts are decoded using a noisy channel model [4] [5] in order to derive a probabilistic transcription (PT) [6], defined as a probability mass function (pmf) over phoneme sequences in the utterance language.

Although PTs are not as accurate as native transcriptions, it has been shown that ASR adapted using PTs is more accurate than cross-lingual or semi-supervised ASR [6].

This paper describes two methods that further improve the error rate of ASR trained using PTs. First, we have previously shown [7] that language specific constraints reduce the error rate of a PT, but [7] only focuses on the text level; this paper demonstrates that the resulting PTs also yield ASR with reduced error. Second, this paper re-imagines probabilistic transcription as a kind of machine translation problem, in which mismatched transcripts are translated

to phone sequences. Building on this analogy, sequence-to-sequence translation models based on recurrent neural networks (RNNs) were trained to generate PTs, from which ASR was trained.

Section 2 briefly describes probabilistic transcriptions, the language-specific constraint method, and the RNN-based PT decoding scheme. Section 3 describes data preparation, along with the experimental details and results, and our main conclusions are listed in Section 4.

## 2. IMPROVING AND CONSTRAINING THE CHANNEL FOR MISMATCHED CROWDSOURCING

A probabilistic transcription (PT) is a distribution over utterance-language phone sequences. When a PT is computed from mismatched crowdsourcing, it can be written as  $\Pr(\pi|T)$ , where  $T$  is a set of mismatched transcripts in the orthography of the annotation language (e.g., English), and  $\pi$  is a phone sequence in the orthography of the utterance language (e.g., Greek). In order to make the computation of  $\Pr(\pi|T)$  tractable, the nonsense texts produced by multiple transcribers are first pruned and merged in order to create a distribution over nonsense texts [4]. Using  $\lambda$  to denote this annotation-language nonsense transcript when viewed as a random variable, the computations can be summarized as:

$$\begin{aligned}\Pr(\pi|T) &= \sum_{\lambda} \Pr(\pi, \lambda|T) \Pr(\lambda|T) \\ &\approx \max_{\lambda} \Pr(\pi|\lambda) \Pr(\lambda|T) \\ &= \max_{\lambda} \left( \frac{\Pr(\lambda|\pi)}{\Pr(\lambda)} \Pr(\pi) \right) \Pr(\lambda|T)\end{aligned}\quad (1)$$

$\Pr(\lambda|\pi)$  is represented by a FST which maps phones to letters [4].  $\Pr(\lambda)$  is a prior over transcriber-language letter sequences,  $\Pr(\lambda|T)$  is a distribution to merge the transcripts in  $T$  over "representative transcripts" denoted by  $\lambda$  and  $\Pr(\pi)$  is a bigram phone language model trained using phone sequences synthesized in the utterance language by applying a character-based grapheme-to-phoneme converter to some utterance-language texts downloaded from Wikipedia. As for the decoding process, the phone sequence which maximizes  $\Pr(\pi|T)$  will be chosen according to the following equation:

$$\pi^* = \arg \max_{\pi} \Pr(\pi|T) \quad (2)$$

### 2.1. Target-Language Post-Processing

A PT is a probability distribution, in a form that can be represented by a weighted finite state transducer (WFST). Language specific

constraints have been applied to reduce the entropy of PTs. Although these constraints are mined from Internet text, the results in [7] showed that they can reduce the error rate of PTs.

Since we have assumed that a pronunciation dictionary is unavailable for the target language, a list of orthographic words in the target language without known pronunciation is used instead. Since available text data in the target language is sparse, we assume every word in this list is equally likely. This language model constraint can be represented as an FST, denoted as **LM**. Although there may not exist a suitable pronunciation model, downloading and appropriately reconstructing the Wikipedia page titled "XXX Alphabet" (where XXX is the test language, for example, Greek in this case), is a possible method to estimate the set of all possible grapheme-to-phoneme (G2P) mappings for the target language. Then for any given grapheme sequence in the test language, the resulting G2P produces all phone sequences (with equal likelihood) that are attested on Wikipedia as possible pronunciations. This kind of unweighted G2P is unlikely to be useful for automatic speech recognition, but is better than having no pronunciation at all: it imposes a loose upper bound on the collection of possible phone sequences that might correspond to any given orthographic sequence. Such G2Ps have been constructed for seventy languages, and are available at [8]. In this paper, we denote the resulting FST mapping graphemes to phonemes as **G2P**.

We regard the FSTs **G2P** and **LM** described above as the available target language information. It is possible to reduce PERs of PTs by only applying this information as constraints on the PTs. The information in these transducers is applied to the **PT** by composing the FSTs, creating the modified probabilistic transcription  $\widehat{\mathbf{PT}} = \mathbf{G2P}^{-1} \circ \mathbf{LM} \circ \mathbf{G2P} \circ \mathbf{PT}$

Furthermore, since the weights on each edge in the G2P and LM are regarded as not useful, it is possible to discount them using zero-exponentiation, defined such that  $x^0 \triangleq 1$  if and only if  $x \neq 1$  but  $0^0 \triangleq 0$ . Using this notation, the edge weights in **G2P** are a model of  $\Pr(W|\pi)^0$  and those in **LM** are a model of  $\Pr(W)^0$ , therefore the edge weights in  $\widehat{\mathbf{PT}}$  are a model of

$$\widehat{PT}(\pi|T) = \sum_W \Pr(\pi|W)^0 \Pr(W)^0 \Pr(W|\pi) \Pr(\pi|T) \quad (3)$$

where  $W$  is the possible word sequence. Eq. (3) shows that **PT** with post processing is a combination of a language-dependent G2P (representing the set of grapheme-to-phoneme transductions that are attested in the target language) and a language-dependent language model LM. As an intermediate step, it is possible to limit PT by simply pruning away phonemes that don't exist in the test language. Phoneme pruning can be computed by the FST composition  $\widehat{\mathbf{PT}} = \mathbf{G2P}^{-1} \circ \mathbf{G2P} \circ \mathbf{PT}$ , which computes:

$$\widehat{PT}(\pi|T) = \sum_W \Pr(\pi|W)^0 \Pr(W|\pi)^0 \Pr(\pi|T) \quad (4)$$

## 2.2. RNN-based Probabilistic Transcriptions

Recurrent neural networks (RNNs) have been successfully applied to a number of sequence prediction tasks. A popular RNN model called RNN Encoder-Decoder was introduced to learn mappings between variable-length sequences which was successfully evaluated on a machine translation task [9]. We observe that the problem of building probabilistic transcriptions could be viewed as a simpler variant of the translation problem in which mismatched transcripts are translated into phone sequences.

The RNN Encoder-Decoder architecture consists of a pair of RNNs – an encoder and a decoder RNN. The encoder RNN processes an entire input sequence sequentially and produces a fixed-length representation after reading the last symbol of the input sequence. The decoder RNN then decodes the given fixed-length vector into a target-language phone sequence. During the training process, the probability of a source sentence given a target sentence is maximized by jointly training the two RNNs.

Formally, we denote the source symbol sequence, as  $X = \{x_1, x_2, \dots, x_T\}$ , and the target sequence as  $Y = \{y_1, y_2, \dots, y_{\hat{T}}\}$  (the lengths of  $X$  and  $Y$  could be different). Following notation introduced in [9], the encoder reads each input symbol into an RNN in which the hidden state  $h_{<t>}$  is updated as:

$$h_{<t>} = f(h_{<t-1>}, x_t) \quad (5)$$

where function  $f$  is a non-linear activation function.

The encoder produces a fixed-length vector (denoted by  $c$ ) as output after reading the entire input sequence; the last symbol of the input sequence is marked by a special end-of-sequence symbol e.g. " $</s>$ ".  $c$  is an encoded representation of the input sequence that is obtained from the hidden state of the encoder RNN at time step  $T$ . Conditioned on this summary vector  $c$ , the hidden states in the decoder RNN are updated by:

$$h_{<t>} = f(h_{<t-1>}, y_{t-1}, c) \quad (6)$$

where  $h_{<t>}$  is also conditioned on the previous symbol in the target sequence,  $y_{t-1}$ . The decoder RNN produces the target sequence according to the following probability distribution (conditioned on both  $c$  and  $y_{t-1}$ ):

$$\Pr(y_t|y_{t-1}, \dots, y_1, c) = g(h_{<t>}, y_{t-1}, c) \quad (7)$$

Here,  $g$  is the softmax activation function. The encoder and decoder RNNs are trained by choosing network parameters in order to maximize the log-probability of the output sequence given the input sequence. i.e.

$$\theta = \operatorname{argmax}_{\theta} \frac{1}{N} \sum_{n=1}^N \log(\Pr_{\theta}(Y_n|X_n)) \quad (8)$$

where  $N$  is the number of training samples and  $\theta$  is the set of parameters in the model.

In this paper, the encoder and decoder RNNs are implemented using Long Short-Term Memory (LSTM) networks, which can learn relatively longer temporal dependencies than simple RNNs. More experiment details are described in section 3.3.

PTs are computed by decoding mismatched transcripts using a WFST model of non-native speech perception. The non-native perceptual model can also be viewed as a sequence-to-sequence translation of non-native phones into letters in the native orthography. Since this process is analogous to translation, it is possible to train and test algorithms that have been effective in machine translation, such as the encoder-decoder RNN architecture. During training, the source sequence is a set of annotation-language orthographic transcripts (acquired from crowd workers), while the target sequence is an utterance-language phone sequence (acquired from a native speaker of the utterance language). In our target application, it may not be possible to recruit native speakers of the target language, therefore the RNN is trained using transcripts of speech in other languages. For example, to train an RNN that translates English mismatched transcription into Greek phones, we use training data that

includes English mismatched transcription of six other languages: Arabic, Cantonese, Dutch, Hungarian, Mandarin, Swahili, and Urdu. Vocabulary for input is English letters (lower case and capital case), and output is International Phonetic Alphabet (IPA).

After training this model by equation (8), we apply it to translate the input sequence  $X$  to the target language phone sequence  $Y$ . Instead of using phone sequences directly from this model, we added a bi-gram language model mapping the target language into the translation process, which could (1) remove phonemes which never occur in the target language, and (2) improve accuracy:

$$Y^* = \underset{Y}{\operatorname{argmax}} \sum_{t=1}^T [\log \Pr_{\theta}(y_t|y_{t-1}, c) + \log \Pr_{\text{LM}}(y_t|y_{t-1})] \quad (9)$$

where  $T$  is the length of the translated sequence  $Y$  and  $\text{LM}$  is the bigram language model of the target language.

### 3. EXPERIMENTS AND RESULTS

Speech data were extracted for eight languages from SBS audio podcasts [10]: Arabic (ARB), Cantonese (YUE), Dutch (NLD), Greek (ELL), Hungarian (HUN), Mandarin (CMN), Swahili (SWH) and Urdu (URD). Native transcriptions were acquired for 40-60 minutes of speech in each of these languages. Cantonese transcriptions were obtained from I2R, Singapore as part of a collaborative research project, Greek transcriptions were collected from native speakers of Greek on the crowdsourcing platform UpWork [11] and speech in the remaining six languages were transcribed by paid student volunteers at the University of Illinois, Urbana-Champaign who were native speakers of these languages. All the native transcriptions were converted into phonemic sequences using a universal phone set, manually constructed as a subset of the IPA. The eight speech datasets were split into roughly 40 minutes of training data and 10 minutes each for the development and evaluation sets.

Mismatched transcriptions were acquired and decoded using published methods [6]. Ten different crowd workers on Amazon Mechanical Turk [12] [13] were asked to provide a sequence of English nonsense syllables for each audio clip.

**Cross-lingual baseline ASR systems:** For all eight languages, we built standard GMM-HMM based ASR systems trained on speech from all languages other than the language being recognized. These ASR systems were implemented using the Kaldi toolkit [14] and are henceforth referred to as **CL**.

**ASR adaptation:** The cross-lingual baseline ASR systems (trained on data other than the target language) were adapted with probabilistic transcriptions for the target language using maximum a posteriori (MAP) estimation. This adaptation process is detailed further in [6].

#### 3.1. Using Language-Specific Pronunciation Constraints

We demonstrate the utility of language-specific G2P mappings using Greek as our target language. Since there is no standard pronunciation dictionary available for Greek, we resort to two simple techniques to derive pronunciation constraints for Greek, described first in [7].

First, a set of G2P mappings were compiled for Greek based on the description of its orthography [8]. These rule sets can be fairly easily generated for a range of different languages, as shown in [8]. Second, a Greek pronunciation dictionary was constructed using data available from the *Translation as a Service (Taas)* project

[15], which has about 200,000 Greek words with corresponding pronunciations available for some words. These pronunciations use an ASCII-based alphabet which we further map to IPA using a deterministic phone mapping.

Specifically, we defined three kinds of constraints [7].

- We first used the inverse of the Greek G2P rules to map the phone-based transcriptions into Greek word sequences and then map them back to phone sequences for evaluation (computing the (PT) transducer of Eq. (4)). This simple constraint forces the probabilistic transcriptions to be matched with valid Greek words. This is referred to as **“G2P”**.
- For the sake of comparison, we also compute phone error rates using the Greek dictionary (described above) in conjunction with the G2P rules (computing the (PT) transducer of Eq. (3)). Greek words appearing in the dictionary are mapped to their corresponding pronunciations and any remaining out-of-vocabulary words are mapped to phones using the Greek G2P rules. This constraint is referred to as **“G2P + dict”**.
- Finally, we also compute the (PT) transducer after imposing word constraints using a bigram word language model (i.e., obtained by replacing  $\Pr(W)^0$  in Eq. (3) with  $\Pr(W)$  derived from a bigram model). Similar to the phone bigram model, we use the Greek Wikipedia text and the CMU CLMTK toolkit to train a bigram word-level language model. We refer to this constraint as **“WLM”**.

Bigram phone language models for Greek were implemented to represent  $\Pr(\pi)$  in Equation (1). In order to train a phone language model, the Greek word sequences were first mapped to phone sequences using the Greek pronunciation dictionary described earlier, along with applying the Greek G2P rules for any out-of-vocabulary words.

Table 1 shows the ASR adaptation results on the Greek development and evaluation sets after using each of the above-mentioned pronunciation constraints. We observe a consistent reduction in phone error rates after incorporating the G2P constraints. Using PT adaptation with the cross-lingual baseline system significantly improves PER performance with a relative PER reduction of 9%. Adaptation with PTs subjected to the G2P constraints further improves ASR performance, with the WLM constraints providing the largest reduction of 10% PER on the evaluation set.

**Table 1.** PERs on the evaluation and development sets of Greek using various language-specific constraints. CL=cross-lingual baseline. PT=PT adaptation. G2P, dict, WLM: defined in the text.

System	PER% (dev)	PER% (eval)
CL	71.39	69.68
CL + PT	64.72	63.69
CL + PT + G2P	62.07	61.45
CL + PT + G2P + dict	60.24	60.11
CL + PT + WLM	58.96	57.23

#### 3.2. Encoder-decoder RNN-based PTs

For the encoder-decoder RNN-based PTs, the experimental setup was designed to be the same as in [6] in order to enable a fair comparison. Seven languages other than Greek were used to train the RNN, which was then tested using Greek.

**Table 2.** Phone error rates of the best path through the PTs based on RNN sequence-to-sequence models, compared to baseline using an FST sequence-to-sequence model (from [6]). Performance on development set in parentheses.

Language (ISO 639-3)	FST Translation	RNN Translation	% Rel. Reduction
ARB	66.2 (65.8)	60.9 (60.3)	8.0 (8.3)
YUE	67.8 (66.4)	62.4 (62.7)	8.0 (5.6)
NLD	70.9 (68.9)	67.9 (64.8)	4.2 (6.0)
HUN	63.5 (63.7)	60.2 (59.6)	5.2 (6.4)
CMN	69.6 (70.9)	67.2 (66.8)	3.4 (5.8)
SWH	50.3 (47.6)	46.8 (43.6)	7.0 (8.4)
URD	70.5 (67.2)	66.9 (64.6)	5.1 (3.9)

**Table 3.** Phone error rates of ASR adapted to PTs. PTs constructed using either FST-based (numbers from [6]) or RNN-based translation models. Dev set in parentheses.

Language code	Cross-lingual (CL)	CL + PT-FST adaptation	CL + RNN based PT adaptation	% Rel. redn over CL+PT-FST
YUE	68.40 (68.35)	57.20 (56.57)	55.62 (56.01)	2.8 (1.0)
HUN	68.62 (66.90)	56.98 (57.26)	53.73 (54.32)	5.8 (5.1)
CMN	71.30 (68.66)	58.21 (57.85)	55.64 (55.90)	4.4 (3.4)
SWH	63.04 (64.73)	44.31 (48.88)	41.21 (44.66)	7.0 (4.2)

Every input and output symbol was encoded using a one-hot representation. The input symbols are English letters and output symbols are phones in IPA. We also introduced two special symbols marking the start of a sentence, "<s>" and the end of a sentence, "</s>". There was no unknown word symbol since both symbol sets were derived from a fixed vocabulary. The input sequence was reversed in order to improve performance (as suggested in [16]). We used two deep LSTMs as our basic model for the encoder and the decoder, respectively and each of them have two hidden layers. Each hidden layer comprised 100 hidden units. A softmax activation function is used at the network output to generate probability distributions across the output phones. The LSTM parameters were initialized using a uniform distribution between -0.1 and 0.1. Stochastic gradient descent with a fixed learning rate of 0.4 was used to train our LSTM network, with a convergence threshold of  $\epsilon=10^{-7}$ . For each iteration, 128 mismatched transcripts and their aligned native transcripts in the training set were selected for training. After 8 epochs, the learning rate was adjusted to 0.2 for more precise updating. These hyper-parameter settings were determined by fine-tuning on the development sets. After training this model, we used a small amount of target language data (10 minutes) to adapt this model further. The LSTM architecture was implemented using Theano libraries [18] and run on GPUs.

Training data include 11520 mismatched transcripts and their corresponding native phonetic transcriptions in languages other than the target language. The development data include 480 mismatched transcripts and their native phonetic transcriptions in target languages.

For each input symbol in the mismatched transcripts, the RNN models compute a probability distribution over the output phones. This output distribution is implemented using a finite state transducer and composed with a phone bigram language model for each language. The bigram language models are constructed using text documents from Wikipedia mined for each of these languages. The word sequences were converted into phone sequences using G2P/pronunciation models before training bigram phone language models (using the CMU CLMTK toolkit [17]). More details

on building these phone language models are specified in [6].

Table 2 shows the 1-best probabilistic phone transcription error rates using PTs derived from the RNN-based models. The numbers in parentheses refer to the 1-best error rates computed using PTs derived using maximum-likelihood alignments between letter sequences and phone sequences trained using the EM algorithm (as defined in [6]); we will refer to them as PT-ML. We observe that the RNN-based PTs consistently outperform PT-MLs across all seven languages. The largest PT error reductions were observed for Swahili and Arabic.

Table 3 demonstrates that the improvement in quality of RNN-based PTs translates to improvements in ASR performance. We see performance improvements in ASR systems built for four different languages, with the largest improvement observed for Swahili.

#### 4. CONCLUSION

This article describes two methods to reduce phone error rates of probabilistic transcriptions (PTs) built from mismatched crowdsourcing by adding language specific constraints, (i.e. a rudimentary rule-based G2P and a list of orthographic word forms mined from the internet) and applying an RNN Encoder-Decoder architecture. ASR built on multilingual transcriptions is adapted by these two optimized PTs. Compared with simple PT-adapted ASR, the post processing method trains an ASR with up to 8.9% relative PER reduction, while the RNN Encoder-Decoder based PT (without any language specific constraints) provides up to 7.0% relative PER reduction.

#### 5. ACKNOWLEDGEMENTS

This work was funded in part by a grant from the DARPA LORELEI program.



## 6. REFERENCES

- [1] S. Novotney and C. Callison-Burch, "Cheap, fast and good enough: Automatic speech recognition with non-expert transcription," in *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 207–215, Association for Computational Linguistics, 2010.
- [2] M. Eskenazi, G.-A. Levow, H. Meng, G. Parent, and D. Suen-dermann, *Crowdsourcing for speech processing: Applications to data collection, transcription and assessment*. John Wiley & Sons, 2013.
- [3] L. Besacier, E. Barnard, A. Karpov, and T. Schultz, "Automatic speech recognition for under-resourced languages: A survey," *Speech Communication*, vol. 56, pp. 85–100, 2014.
- [4] P. Jyothi and M. Hasegawa-Johnson, "Acquiring speech transcriptions using mismatched crowdsourcing.," in *AAAI*, pp. 1263–1269, 2015.
- [5] P. Jyothi and M. Hasegawa-Johnson, "Transcribing continuous speech using mismatched crowdsourcing," in *Proceedings of Interspeech*, 2015.
- [6] C. Liu, P. Jyothi, H. Tang, V. Manohar, R. Sloan, T. Kekona, M. Hasegawa-Johnson, and S. Khudanpur, "Adapting asr for under-resourced languages using mismatched transcriptions," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5840–5844, IEEE, 2016.
- [7] X. Kong, P. Jyothi, and M. Hasegawa-Johnson, "Performance improvement of probabilistic transcriptions with language-specific constraints," *Procedia Computer Science*, vol. 81, pp. 30–36, 2016.
- [8] M. Hasegawa-Johnson, "Ws15 dictionary data," *Downloaded 11/15/2015 from <http://isle.illinois.edu/sst/data/dict>*, 2015.
- [9] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," *arXiv preprint arXiv:1406.1078*, 2014.
- [10] "Special broadcasting services australia," *Downloaded 1/26/2016 from, <http://www.sbs.com.au/podcast/yourlanguage>*.
- [11] "Upwork downloaded 1/26/2016 from <https://www.upwork.com>,"
- [12] "Amazon Mechanical Turk, downloaded 1/26/2016 from <http://www.mturk.com>,"
- [13] E. Pavlick, M. Post, A. Irvine, D. Kachae, and C. Callison-Burch, "The language demographics of amazon mechanical turk," *Transactions of the Association for Computational Linguistics*, vol. 2, pp. 79–92, 2014.
- [14] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, et al., "The kaldi speech recognition toolkit," in *IEEE 2011 workshop on automatic speech recognition and understanding*, no. EPFL-CONF-192584, IEEE Signal Processing Society, 2011.
- [15] "Translation as a Service (TASS), downloaded 1/26/2016 from <http://www.taas-project.eu/>,"
- [16] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Advances in neural information processing systems*, pp. 3104–3112, 2014.
- [17] "Carnegie Mellon University Language Modeling Toolkit CMUCLMTK, downloaded 1/26/2016 from <http://cmusphinx.sourceforge.net/wiki/>,"